arXiv:1610.03524v1 [cs.DC] 11 Oct 2016

Improved Parallel Construction of Wavelet Trees and Rank/Select Structures

Julian Shun*
University of California, Berkeley
jshun@eecs.berkeley.edu

Abstract

Existing parallel algorithms for wavelet tree construction have a work complexity of $O(n \log \sigma)$. This paper presents parallel algorithms for the problem with improved work complexity. Our first algorithm is based on parallel integer sorting and has either $O(n\sqrt{\log\log n}\lceil \log\sigma/\sqrt{\log n}\rceil)$ work and polylogarithmic depth, or $O(n\lceil \log\sigma/\sqrt{\log n}\rceil)$ work and sub-linear depth. We also describe another algorithm that has $O(n\lceil \log\sigma/\sqrt{\log n}\rceil)$ work and $O(\sigma + \log n)$ depth. We then show how to use similar ideas to construct variants of wavelet trees (arbitrary-shaped binary trees and multiary trees) as well as wavelet matrices in parallel with lower work complexity than prior algorithms. Finally, we show that the rank and select structures on binary sequences and multiary sequences, which are stored on wavelet tree nodes, can be constructed in parallel with improved work bounds. In particular, we show that the rank and select structures can be constructed for binary sequences in $O(n/\log n)$ work and $O(\log n)$ depth, and for multiary sequences in $O(n \log\sigma/\log n)$ work and $O(\log n)$ depth. These work bounds match that of the best existing sequential algorithms for constructing rank and select structures.

1 Introduction

The *wavelet tree* is a space-efficient data structure that supports access, rank, and select queries on a sequence in time logarithmic in the alphabet size. It was introduced by Grossi et al. [12], who used it to design a compressed suffix array. Wavelet trees have many other applications [19, 17]—for example, they can be used to obtain compressed representations of sequences, permutations, grids, graphs, and self-indexes based on the Burrows-Wheeler transform, and can also be used for two-dimensional range queries [16].

The standard sequential method for constructing a wavelet tree on a sequence of length n with alphabet size σ takes $O(n \log \sigma)$ work. Very recently, faster sequential algorithms with $O(n\lceil \log\sigma/\sqrt{\log \sigma}\rceil)$ work have been described [18, 1]. As for prior parallel algorithms, Fuentes-Sepulveda et al. [11] present algorithms that require $O(n \log \sigma)$ work and $O(n)$ depth (parallel time). Shun [22] improved the result by developing faster parallel algorithms with $O(n \log \sigma)$ work and polylogarithmic depth. Recently, Labeit et al. [15] present more space-efficient parallel wavelet tree algorithms with the same work and depth bounds as [22].

*Supported by the Miller Institute for Basic Research in Science at UC Berkeley.

While parallel algorithms exist for wavelet tree construction, their work complexity is higher than that of the best sequential algorithms, which take $O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$ work [18, 1]. This paper presents parallel algorithms for wavelet tree construction with improved work complexity. Our first algorithm is a parallelization of the algorithm in [1] and uses parallel integer sorting. Depending on the parallel integer sorting subroutine used, the algorithm takes either $O(n\sqrt{\log\log n}\lceil\log\sigma/\sqrt{\log n}\rceil)$ work and $O(\log n\log\sigma)$ depth or $O((n/\epsilon)\lceil\log\sigma/\sqrt{\log n}\rceil))$ work and $O(\log\sigma(\log n+(n^\epsilon/(\epsilon\sqrt{\log n}))))$ depth for a constant $0<\epsilon<1$. This results in either a polylogarithmic-depth algorithm with improved work complexity, or a sub-linear depth algorithm whose work matches that of the best sequential algorithm. Our second algorithm is based on a simple domain-decomposition approach as used in [11, 15], and takes $O(\sigma P+n\lceil\log\sigma/\sqrt{\log n}\rceil)$ work and $O((n/P)\lceil\log\sigma/\sqrt{\log n}\rceil+\log P)$ depth for a parameter $P\geq 1$. Setting $P=\Theta((n/\sigma)\lceil\log\sigma/\sqrt{\log n}\rceil)$ gives an algorithm with $O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$ work and $O(\sigma+\log P)$ depth. This algorithm therefore has high parallelism for small alphabet sizes.

Using similar ideas we also obtain improved algorithms for constructing variants of the standard wavelet tree, such as arbitrary-shaped binary wavelet trees [10], multiary trees [9], and wavelet matrices [7]. Wavelet tree nodes store rank and select structures, and so to achieve the improved work bounds, we show how to construct in parallel the rank and select structures of binary and multiary sequences work-efficiently. For binary sequences of length n we show how to construct the structures in $O(n/\log n)$ work and $O(\log n)$ depth (the sequence lengths across all wavelet tree nodes sum to $O(n\log\sigma)$, so this contributes a total of $O(n\log\sigma/\log n)$ work). For sequences containing characters in $[0,\ldots,\sigma-1]$ we show how to construct the structures in $O(n\log\sigma/\log n)$ work and $O(\log n)$ depth. The work bounds match those of the sequential algorithms described in [1]. This is the most technically involved part of the paper and obtaining these bounds in parallel requires carefully packing values into words, working on the compact representations, constructing appropriate lookup tables, and defining appropriate operators for prefix sum computations. Existing and new bounds for the problems studied in this paper are shown in Table 1.

2 Preliminaries

We analyze algorithms in the *work-depth model*, where the *work* W is the number of operations required (equivalent to the standard sequential time complexity) and the *depth (parallel time)* D is the length of the longest critical path in the computation [14]. The parallelism of an algorithm is equal to W/D. With p available processors, using Brent's scheduling theorem [4] we can bound the running time by $W/p+D$. Algorithms with low work and depth complexities can be executed efficiently on parallel machines using a work-stealing scheduler [3]. We say that a parallel algorithm is *work-efficient* if its asymptotic work complexity matches that of the best sequential algorithm. As in the standard word RAM model, we assume that $\Theta(\log n)$ bits fit in a word, and reading or writing a word requires unit work.

A sequence of symbols will be denoted by S, its length by n, and its alphabet size by σ. For a sequence S, $\mathsf{access}(S,i)$ returns the symbol at position i of S, $\mathsf{rank}_c(S,i)$ returns the number of times c appears in S from positions 0 to i, and $\mathsf{select}_c(S,i)$ returns the position of the i'th occurrence of c in S. A *wavelet tree* is a data structure supporting access, rank, and select operations on a sequence in $O(\log\sigma)$ work [12].[1] A standard wavelet tree is a balanced binary tree where each node represents a range of the symbols in the alphabet using a bitmap (binary sequence). We assume that $\sigma\leq n$, and that the alphabet is $[0,\ldots,\sigma-1]$, as the symbols can be mapped to a contiguous range otherwise. The structure of the wavelet tree is defined recursively as follows: The root represents the symbols $[0,\ldots,2^{\lceil\log\sigma\rceil}-1]$. A node v which represents the symbols $[a,\ldots,b]$ stores a bitmap

[1] We use $\log x$ to mean the base 2 logarithm of x unless specified otherwise.

Data Structure	Algorithm	Work	Depth
Binary Wavelet Tree	Sequential [1, 18]	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	–
	[22, 15]	$O(n\lceil\log\sigma\rceil)$	$O(\log n\log\sigma)$
	This paper	$O(n\sqrt{\log\log n}\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\log n\log\sigma)$
	This paper	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\log\sigma(\log n+n^\epsilon/\sqrt{\log n}))$
	This paper	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\sigma+\log n)$
Arbitrary-shaped Binary Wavelet Tree (height h)	Sequential [1, 18]	$O(n\lceil h/\sqrt{\log n}\rceil)$	–
	[22]	$O(nh)$	$O(h\log n)$
	This paper	$O(n\sqrt{\log\log n}\lceil h/\sqrt{\log n}\rceil)$	$O(h\log n)$
	This paper	$O(n\lceil h/\sqrt{\log n}\rceil)$	$O(h(\log n+n^\epsilon/\sqrt{\log n}))$
Multiary Wavelet Tree (degree $d=o(\log^{1/3}n)$)	Sequential [1, 18]	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	–
	[22]	$O(n\log\sigma)$	$O(\log n\log\sigma)$
	This paper	$O(n\sqrt{\log\log n}\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\log n\log\sigma)$
	This paper	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\log\sigma(\log n+n^\epsilon/\sqrt{\log n}))$
Wavelet Matrix	[22]	$O(n\log\sigma)$	$O(\log n\log\sigma)$
	This paper	$O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$	$O(\log n\log\sigma)$
Binary Rank and Select	Sequential [1, 18]	$O(n/\log n)$	–
	[22]	$O(n)$	$O(\log n)$
	This paper	$O(n/\log n)$	$O(\log n)$
Generalized Rank and Select	Sequential [1]	$O(n\log\sigma/\log n)$	–
	[22]	$O(n)$	$O(\log n)$
	This paper	$O(n\log\sigma/\log n)$	$O(\log n)$

Table 1: New and existing work and depth bounds for constructing data structures. We omit the depth term for the sequential algorithms.

which has a 0 in position i if the i'th symbol in the range $[a,\ldots,b]$ is in $[a,\ldots,(a+b+1)/2-1]$ and 1 otherwise. It will have a left child that represents the symbols $[a,\ldots,(a+b+1)/2-1]$ and a right child that represents the symbols $[(a+b+1)/2,\ldots,b]$. The recursion stops when the range is of size 2 or less or if a node has no symbols to represent. An example of a wavelet tree is shown in Figure 1. We point out that the original wavelet tree description in [12] uses a root whose range is not necessarily a power of 2, but the definition here gives the same asymptotic query times and leads to a simpler description of our construction algorithms.

Each node in the wavelet tree stores a *succinct* rank/select structure on its bitmap (whose size is sub-linear in the bitmap length) to allow for constant-work rank and select queries. The bitmaps and the rank/select structures together take $n\lceil\log\sigma\rceil+o(n\log\sigma)$ bits of space. The wavelet tree topology requires $O(\sigma\log n)$ bits to store pointers, but this can be reduced or removed by modifying how the queries are performed [16, 6].

Our algorithms use prefix sum as a parallel primitive [14]. *Prefix sum* takes as input an array X of length n, an associative binary operator $\oplus$, and an identity element $\perp$ such that $\perp\oplus x=x$ for any x, and returns the array $(\perp,\perp\oplus X[0],\perp\oplus X[0]\oplus X[1],\ldots,\perp\oplus X[0]\oplus X[1]\oplus\ldots\oplus X[n-2])$, as well as the overall sum $\perp\oplus X[0]\oplus X[1]\oplus\ldots\oplus X[n-1]$. Assuming that applying the $\oplus$ takes constant work, prefix sums can be implemented in $O(n)$ work and $O(\log n)$ depth [14]. Unless specified otherwise, we will use $\oplus$ to be the addition operator on integers.

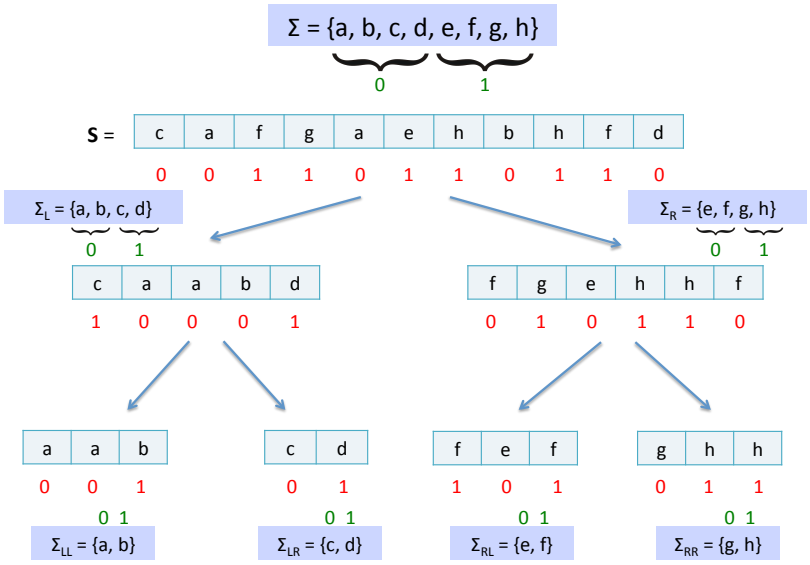
Σ = {a, b, c, d, e, f, g, h}
0 1
S = c a f g a e h b h f d
0 0 1 1 0 1 1 0 1 1 0
Σ_L = {a, b, c, d}
0 1
Σ_R = {e, f, g, h}
0 1
c a a b d
1 0 0 0 1
f g e h h f
0 1 0 1 1 0
a a b
0 0 1
0 1
Σ_LL = {a, b}
c d
0 1
0 1
Σ_LR = {c, d}
f e f
1 0 1
0 1
Σ_RL = {e, f}
g h h
0 1 1
0 1
Σ_RR = {g, h}

Figure 1: (Best viewed in color.) Example of a wavelet tree on the sequence $\mathsf{S} = cafgaehbhfd$ *and where* $\sigma = 8$ *and the alphabet is* $\Sigma = \{a, b, c, d, e, f, g, h\}$*. The root contains the entire sequence as well as a bitmap (shown in red) indicating whether the symbol comes from the left or the right half of* Σ*. The two children* L *and* R *of the root contain the characters in the left and right half of* Σ*, respectively. The bitmap for the left (right) child is constructed by checking if the character comes from the left half or the right half of* Σ_L *(*Σ_R*). The leaves each represent an alphabet of size* 2*.*

3 Review of the $O(n\lceil \log \sigma / \sqrt{\log n} \rceil)$ Work Sequential Algorithm

We first review how the $O(n\lceil \log \sigma / \sqrt{\log n} \rceil)$ work sequential wavelet tree construction algorithm from [1] works, as we will be parallelizing this algorithm. A similar sequential algorithm was independently described in [18]. The main idea of the algorithm is to process integers of size smaller than $\log \sigma$. The basic data structure used is a *packed list*, which stores N b-bit integers using $\lceil Nb/\log n \rceil$ words. It supports appending a length N list in $O(1 + \lceil Nb/\log n \rceil)$ work and splitting a list into smaller lists of at most length k in $O(\lceil Nb/\log n \rceil + N/k)$ work. A *big node* is defined to be a node at a distance that is a multiple of τ from the root, where τ is a parameter to be chosen. Big nodes store the elements that it represents in S, using $\log \sigma$ bits per element as in the standard representation. These can be computed recursively as follows. The root stores S. Assume that the sub-sequences for the big nodes at distance $\alpha\tau$ are already computed. Then to compute the elements in the big nodes at distance $(\alpha + 1)\tau$, the big nodes at distance $\alpha\tau$ look at the τ bits starting at position $\alpha\tau$ in the binary representation of each element to determine which of its descendant big nodes at distance $(\alpha + 1)\tau$ to place the element at (there are 2^τ such descendants). The computation for big nodes requires $O(n\lceil \log \sigma/\tau \rceil)$ work.

Nodes at all other levels of the tree only need to store τ bits per element (the τ bits starting at position $\alpha\tau$ where $\alpha\tau$ is the level of its nearest big node ancestor) because there are only τ levels between two big node levels. These nodes use *short lists* to store τ-bit integers containing the τ relevant bits of the elements it represents. These are stored as packed lists. Computing the bitmap values and short lists is done recursively. The short lists of the children of a big node can be computed by extracting the relevant bits from the elements of the big nodes in $O(n\lceil \log \sigma/\tau \rceil)$

work across all big nodes. Given a short list of a node, computing its own bitmap values and the short lists of its children is done via table lookup. For all packed lists L of at most $\log n/(2\tau)$ τ-bit integers, the bitmap value, and the packed lists L_0 and L_1 consisting of the elements of L whose t'th most significant bit is 0 or 1, respectively are pre-computed for all $t \in [0, \tau - 1]$. Pre-computing this table involves evaluating all $O(2^{\log n/(2\tau)})$ τ-bit integer sequences of length at most $\log n/(2\tau)$ for each value of t. This can be done in $O(n)$ work. Each node splits its short list into blocks of length at most $\log n/(2\tau)$, performs table lookups for each block, and then appends the resulting bitmap values together, L_0's together and L_1's together. The bitmap values are stored in the bitmap associated with the current node, and L_0 and L_1 are passed to its children. For a node with a short list of length N, the total work required is $O(N\tau/\log n)$ as the splitting and merging can be done in $O(N\tau/\log n)$ work overall and table lookups in constant work per block. The lengths of all short lists is $O(n \log \sigma)$, so the total work required for this computation is $O(n\lceil \log \sigma(\tau/\log n)\rceil)$.

The overall work is $O(n\lceil \log \sigma/\tau\rceil + n\lceil \log \sigma(\tau/\log n)\rceil)$ and choosing $\tau = \sqrt{\log n}$ minimizes the work, giving a bound of $O(n\lceil \log \sigma/\sqrt{\log n}\rceil)$.

4 Parallel Wavelet Tree Algorithms

This section first describes how to parallelize the algorithm of [1]. Then we present a simple domain-decomposition based parallel construction algorithm that is work-efficient and whose parallelism depends linearly on σ, and has low depth for small alphabets.

Parallelizing the algorithm of [1]. The sub-sequences for the big nodes can be computed level-by-level using parallel integer sorting. In particular, given the correct sub-sequence S' for a big node at level $\alpha\tau$, we compute the sub-sequences for its big node descendants at level $(\alpha + 1)\tau$ by performing an integer sort on S', where the key for the sort is the value of the (up to) τ bits starting from the $\alpha\tau$'th highest bit of the symbol.

The parallel integer sort that we use is required to be stable since we need to keep the relative ordering among the characters in each descendant node. Unfortunately the only known method for stable parallel integer sorting in linear work and polylogarithmic depth [20] requires the range of the keys of the values being sorted to be polylogarithmic, which does not hold for the value of τ that we will choose. Instead we can either use an algorithm that is not work-efficient, requiring $O(n \log \log n)$ work and $O(\log n)$ depth [21, 2],[2] or use a work-efficient algorithm with $O(n/\epsilon)$ work and $O(n^\epsilon/\epsilon)$ depth for a constant $0 < \epsilon < 1$ [23]. This gives an overall complexity for constructing big nodes of either (a) $O(n \log \log n\lceil \log \sigma/\tau\rceil)$ work and $O(\log n\lceil \log \sigma/\tau\rceil)$ depth or (b) $O((n/\epsilon)\lceil \log \sigma/\tau\rceil)$ work and $O((n^\epsilon/\epsilon)\lceil \log \sigma/\tau\rceil)$ depth for constructing the big nodes.

The lookup table for computing short lists can be pre-computed by evaluating all $O(2^{\log n/(2\tau)})$ τ-bit integer sequences of length at most $\log n/(2\tau)$ for each $t \in [0, \tau - 1]$ in parallel, and storing the answer for each in a unique location.[3] Each sequence and value of t is evaluated sequentially. Overall, this requires $O(\log n)$ depth and $o(n)$ work.

Computing short lists for children of a big node can be done in linear work and $O(\log n)$ depth by extracting the relevant bits from the symbols in the big node, performing prefix sums to get the appropriate offsets, and copying the τ bits of a symbol into the appropriate location in an array of the appropriate child in parallel. Groups of τ-bit integers that together form a word are then packed together and copied into one entry of the short list for the corresponding child in parallel. The bitmaps of the children of a big node can be computed in linear work and $O(\log n)$ depth simply

[2] These algorithms either use randomization or require super-linear space.

[3] For example, initialize an array of size $2^{\log n/(2\tau)} \log n/2$ and for each possible sequence and value of t, use the binary value of the sequence multiplied by $\log n/(2\tau)$ plus the number of leading zeros as its offset into the array.

by extracting the relevant bit from the symbols and packing them together. Computing short lists of other nodes requires merging and splitting packed lists. For each short list, we split it into chunks of length at most $\log n/(2\tau)$ τ-bit integers by copying the relevant bits of each chunk into its own word in constant depth. The algorithm performs a table lookup for each chunk to obtain the parts of the packed lists L_0 and L_1 that the chunk contributes to as well as the part of the bitmap associated with the chunk. All table lookups are done in parallel in constant depth. We then merge together the results to form each of L_0, L_1, and the bitmap for the node. To merge the results of one of the lists together, we compute the length (in bits) of the result associated with each chunk, perform a prefix sum to determine the total length (in bits) and also the offset for each result in a new array, and allocate a new array of the desired length. We then identify the groups of chunks that will copy into the same word, again using prefix sums (some chunks will copy into two words, but this only increases the work by a constant factor). Then, in parallel, all groups merge their chunks sequentially using the packed list operations described in Section 3 and then copy their word into the new array at the appropriate offset. There are a total of $\lceil 2N\tau/\log n\rceil$ chunks if the short list contains N integers, each of which generates a partial result for L_0, L_1, and the bitmap, and so the prefix sum and copying takes $O(N\tau/\log n)$ work and $O(\log n)$ depth (there is a constant-factor overhead due to some chunks not being full, however the complexity is not affected). The overall work for computing the short lists is $O(n\lceil\log\sigma(\tau/\log n)\rceil)$ as in the sequential algorithm. The depth is $O(\log n\log\sigma)$ as there are $\log\sigma$ levels, each requiring $O(\log n)$ depth.

To minimize the overall work we set $\tau = \sqrt{\log n\log\log n}$ when using the $O(n\log\log n)$ work integer sort and $\tau = \sqrt{\log n}$ when using the $O(n/\epsilon)$ work integer sort. Assuming that constructing the rank and select data structures per node can be done in the same bounds, which we describe in Section 5, we obtain the following theorem:

Theorem 4.1. *Wavelet tree construction can be performed in $O(n\sqrt{\log\log n}\lceil\log\sigma/\sqrt{\log n}\rceil)$ work and $O(\log n\log\sigma)$ depth or $O((n/\epsilon)\lceil\log\sigma/\sqrt{\log n}\rceil)$ work and $O(\log\sigma(\log n + n^\epsilon/(\epsilon\sqrt{\log n})))$ depth for a constant $0 < \epsilon < 1$.*

Note that both parallel algorithms described above improve upon the $O(n\log\sigma)$ work complexity of the algorithms described in [22, 15]. The first algorithm has polylogarithmic depth but does not achieve the $O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$ work bound of the best sequential algorithm, while the second algorithm is work-efficient with sub-linear (but not polylogarithmic) depth. However, as long as the number of processors is sub-linear, the second algorithm can make full use of all of the available processors (recall Brent's scheduling theorem and the definition of parallelism from Section 2). Improving parallel integer sorting algorithms would immediately improve the complexity of the wavelet tree algorithms.

Domain-decomposition approach. Another way to construct the wavelet tree in parallel is to use a domain-decomposition approach as done in [11] and [15]. For a parameter P, this approach first splits the input sequence into P equal-sized sub-sequences, constructs the wavelet tree across all sub-sequences in parallel using a sequential algorithm for each, and then merges the bitstrings on the nodes of the P trees together. Constructing the tree for each sub-sequence can be done by using an $O(n\lceil\log\sigma/\sqrt{\log n}\rceil)$ work sequential algorithm [18, 1] in a black-box fashion (where the alphabet size for each sub-sequence is treated as the same as the alphabet size of the entire sequence). The overall work for this step is $O(\sigma P + n\lceil\log\sigma/\sqrt{\log n}\rceil)$ and depth is $O((n/P)\lceil\log\sigma/\sqrt{\log n}\rceil)$.

To merge together the bitstrings, we first form the structure of the final wavelet tree, which takes $O(\sigma)$ work and $O(1)$ depth. Following the idea described in [15], for each node in the final tree structure, we then perform a prefix sums across the lengths of the bitstrings on the corresponding nodes in the P sub-problems (the length is 0 if the node does not exist). This gives the length

of the bitstring on the node in the final tree as well as an appropriate offset into the bitstring for each sub-problem. Then each sub-sequence copies its bitstring into the bitstring of the node in the final tree in parallel at word granularity. The words where multiple sub-sequences can copy into in parallel are marked beforehand to avoid conflicts and handled specially (these "boundary" words can be identified by looking at the offsets of the $O(\sigma P)$ nodes, and there can be at most $O(\sigma P)$ of them). Summed over all nodes in the final tree, the prefix sums take $O(\sigma P)$ work and $O(\log P)$ depth. The extra space required is $O(\sigma P)$. Excluding the special words, the copying takes $O(\sigma P + n\lceil \log \sigma / \log n \rceil)$ work and $O(1)$ depth in total (the σP term in the work comes from the fact that each sub-sequence must spend at least $O(1)$ work per node, and the $\log n$ in the denominator of the second term of the work is because we are copying at word granularity). The special words can all be computed in parallel, taking $O(\sigma P)$ work and $O(\log P)$ depth by concatenating the up to P bitstrings for each special word in a binary fashion. This gives the following theorem:

Theorem 4.2. *A wavelet tree can be constructed in $O(\sigma P + n\lceil \log \sigma / \sqrt{\log n} \rceil)$ work and $O((n/P)\lceil \log \sigma / \sqrt{\log n} \rceil + \log P)$ depth for an integer $P \geq 1$.*

The domain-decomposition algorithm is work-efficient if $P = O((n/\sigma)\lceil \log \sigma / \sqrt{\log n} \rceil)$. Setting $P = \Theta((n/\sigma)\lceil \log \sigma / \sqrt{\log n} \rceil)$ gives the maximum parallelism while achieving work-efficiency, and gives a depth of $O(\sigma + \log P)$. Thus this algorithm has good parallelism when σ is small. To keep the space $O(n)$, the maximum parallelism is achieved when $P = \Theta(n/\sigma)$, giving $O(\sigma\lceil \log \sigma / \sqrt{\log n} \rceil + \log P)$ depth.

4.1 Variants

Arbitrarily-shaped binary trees. The wavelet tree construction algorithm can be extended to binary trees of other shapes (e.g., Huffman-shaped wavelet trees [10]) if the tree structure can be computed efficiently. In particular, the algorithm needs a codeword for each symbol determined by the path from the root to the node representing the symbol in the tree. The codeword is a bitstring, where the i'th most significant bit is 0 if the $(i+1)$'st node in the path is a left child of the i'th node in the path, and is 1 otherwise. Let h be the height of the tree. We assume a lookup table storing a mapping from codeword to symbol. Since the codewords are of length $O(\log h) = O(\log n)$ (since $h \leq \sigma$ and we assumed $\sigma \leq n$), we can access the codeword in constant-work, and construct the lookup table in $O(\sigma)$ work and $O(\log n)$ depth. (We note that codewords for a Huffman-shaped wavelet tree can be generated in $O(n)$ work and $O(\sigma + \log n)$ depth [8, 22].)

The algorithm proceeds as before, where big nodes are constructed every τ'th level in the tree by using integer sorting on τ bits. Some of the combinations of the bits may not correspond to a symbol, and no big nodes are generated for those combinations. The complexity per level is equal to the complexity of integer sorting, and summing across all h/τ levels gives the following bounds for constructing big nodes: (a) $O(n\sqrt{\log\log n}\lceil h/\sqrt{\log n} \rceil)$ work and $O(h\sqrt{\log n} + \log n)$ depth (by setting $\tau = \sqrt{\log n \log\log n}$) or (b) $O((n/\epsilon)\lceil h/\sqrt{\log n} \rceil)$ work and $O(hn^{\epsilon}/(\epsilon\sqrt{\log n}))$ depth for $0 < \epsilon < 1$ (by setting $\tau = \sqrt{\log n}$). The remaining nodes are computed using short lists as before, and the overall work for these nodes is $O(n\lceil h(\tau/\log n) \rceil)$ and depth is $O(h \log n)$. This gives the following theorem, whose work bound improves upon the parallel construction described in [22]:

Theorem 4.3. *Given codewords for the symbols, a binary wavelet tree of height h can be constructed in $O(n\sqrt{\log\log n}\lceil h/\sqrt{\log n} \rceil)$ work and $O(h \log n)$ depth or $O((n/\epsilon)\lceil h/\sqrt{\log n} \rceil)$ work and $O(h(\log n + n^{\epsilon}/(\epsilon\sqrt{\log n})))$ depth for a constant $0 < \epsilon < 1$.*

Multiary wavelet trees. We now describe how to extend the algorithm to construct multiary wavelet trees [9] of degree $d = o(\log^{1/3} n)$ where d is a power of two.[4] Each node now has d children and the sequence that a node stores contains values in $[0, \ldots, d-1]$ instead of being binary as in the standard wavelet tree. We describe the algorithm for balanced trees but the ideas also apply to trees of arbitrary shapes as long as the codewords are provided as input. Similar to the approach of [18] we generate the full binary tree, but only keep sequences for the nodes at levels $\beta \log d$ in the full binary tree for $\beta = [0, \ldots, \log \sigma)$. Each node with a sequence that is kept belongs to the multiary wavelet tree, and if it is at level $\beta \log d$ in the binary tree, its d children are at level $(\beta + 1) \log d$ in the binary tree. With an appropriate numbering scheme (i.e., the children of node i are stored at locations $2i + 1$ and $2i + 2$), the d children of a node can be identified in $O(d)$ work and $O(1)$ depth, contributing $O(\sigma)$ work and $O(1)$ depth overall. The bound from Theorem 4.1 then applies, giving the following theorem which improves upon the $O(n \log \sigma)$ work of the parallel algorithm for multiary wavelet trees from [22].

Theorem 4.4. *A multiary wavelet tree of degree $d = o(\log^{1/3} n)$ where d is a power of two can be constructed in $O(n\sqrt{\log \log n}\lceil \log \sigma / \sqrt{\log n} \rceil)$ work and $O(\log n \log \sigma)$ depth or $O((n/\epsilon)\lceil \log \sigma / \sqrt{\log n} \rceil)$ work and $O(\log \sigma (\log n + n^{\epsilon}/(\epsilon\sqrt{\log n})))$ depth for $0 < \epsilon < 1$.*

We note that each node of a multiary wavelet tree requires storing a generalized rank and select structure on its string, and we describe how to construct the structures within the bounds of Theorem 4.4 in Section 5.2. Babenko et al. [1] show that their multiary wavelet tree algorithm can be used to support range rank and select queries[5] in $O(\log n / \log \log n)$ time, and we point out that we can use our parallel algorithm to improve the efficiency of the pre-processing step. In addition to generating the multiary wavelet tree, Babenko et al. [1] perform some additional pre-processing on the strings of each node, which can be parallelized using the same ideas that we will discuss in Section 5.2.

Wavelet matrix. The wavelet matrix [7] is a variant of the wavelet tree where for level l, all symbols with a 0 as their l'th highest bit are represented on the left side of the level's sequence and all symbols with a 1 as their l'th highest bit are represented on the right side. The relative ordering among the symbols from the previous level is preserved. The wavelet matrix has $O(\log \sigma)$ levels. An $O(n \log \sigma)$ work, polylogarithmic depth parallel algorithm for constructing the wavelet matrix was described in [22]. In this section, we describe how to reduce the work complexity to $O(n\lceil \log \sigma / \sqrt{\log n} \rceil)$.

We will process the bits of the symbols in chunks of $\sqrt{\log n}$ bits and construct the matrix level-by-level. On level $\alpha\sqrt{\log n}$ for each $\alpha \in [0, \ldots, \lceil \log \sigma / \sqrt{\log n} \rceil)$, we extract the $\sqrt{\log n}$ bits of each symbol starting at the $\alpha\sqrt{\log n}$'th position, except for $\alpha = \lceil \log \sigma / \sqrt{\log n} \rceil - 1$, where we simply extract the remaining bits. Constructing levels $\alpha\sqrt{\log n}$ to $(\alpha + 1)\sqrt{\log n} - 1$ of the wavelet matrix will require only these bits of the symbols. We group $\sqrt{\log n}/2$ $\sqrt{\log n}$-bit entries into a chunk, packing them into a word, and use the packed list representation as in the wavelet tree algorithm. We use a lookup table storing all possible bitstrings of up to length $\log n/2$, which for each chunk determines which symbols go to the left and which go to the right, as well as the bitmap, in $O(1)$ work. The lookup table can be computed in $O(\log n)$ depth and $o(n)$ work. Similar to the wavelet tree algorithm, each chunk can be split into two parts, the first that goes to the left side of the sequence and the second that goes to the right. Prefix sums and grouping of chunks are then used

[4] The restriction $d = o(\log^{1/3} n)$ is due to the rank and select structures from [1] that we parallelize.

[5] The *range select* query takes as input an array A of integers in $[1, \ldots, n]$ and integers i, j, k, and returns the k'th smallest element in $A[i, \ldots, j]$. The *range rank* query takes the same inputs and returns the number of elements in $A[i, \ldots, j]$ that are smaller than $A[k]$.

on the packed lists to create the bitmap for the current level as well as the sequence at the next level. On each level, this takes $O(n/\sqrt{\log n})$ work and $O(\log n)$ depth. Summing over all levels gives $O(n\lceil \log\sigma/\sqrt{\log n}\rceil)$ work and $O(\log n \log\sigma)$ depth. This gives the following theorem:

Theorem 4.5. *Wavelet matrix construction can be performed in* $O(n\lceil \log\sigma/\sqrt{\log n}\rceil)$ *work and* $O(\log n \log\sigma)$ *depth.*

Note that the bound is stronger than that of Theorem 4.1 because the algorithm does not use integer sorting. The wavelet tree algorithm uses integer sorting to determine which symbols belong to which tree nodes when computing "big" nodes, whereas the wavelet matrix does not have any tree nodes and does not need this step. Simply extracting the $\sqrt{\log n}$ relevant bits every $\sqrt{\log n}$'th level is sufficient for obtaining the improved work complexity.

5 Improved Parallel Construction of Rank/Select Structures

Wavelet trees and matrices require each node to store a succinct rank and select structure on its bitmaps (or sequences of $(\log d)$-bit values). We show how to construct these structures in parallel within the bounds of the construction algorithms described above.

5.1 Binary Sequences

We first describe the binary sequence case. The goal is to construct the rank/select structures on n bits in $O(n/\log n)$ work to match the work bound of the sequential construction algorithms in [1]. The overall work for rank/select construction in a wavelet tree will therefore be $O(n\log\sigma/\log n)$, which is within the work bound of our parallel algorithms. We assume that the bit sequence is packed into $n/\log n$ words, which is provided by our wavelet tree algorithms from Section 4.

Rank. For rank queries, we use the structure of Jacobson [13]. The data structure divides the bit sequence into *ranges* of size $\log^2 n$. It computes the rank (both the number of 0's before and the number of 1's before) for the last bit in each range. The ranges are further divided into *sub-ranges* of size $\log n$, where the rank of every $\log n$'th bit relative to the beginning of the range is stored. Inside a sub-range, the rank of a position relative to the beginning of the sub-range can be answered with at most two table lookups, where the table stores the answers to all queries of sequences of up to length $\log n/2$.

We initialize two arrays, A_0 and A_1, of length $n/\log n$ each, and for each of the $n/\log n$ words, we count the number of 0's and 1's in the word and store them into its position in the appropriate array. Counting the number of 0's and 1's in a word can be done in $O(1)$ work using the same lookup table as for answering rank queries. Then we compute the prefix sums over each of A_0 and A_1. Then, every $\log^2 n$'th entry in A_0 (A_1) gives the rank of 0 (1) for the last position in each range. The results for the sub-ranges are computed by taking each remaining entries in A_0 and A_1, and subtracting the rank stored for the closest range to the left. The prefix sums require $O(n/\log n)$ work and $O(\log n)$ depth. The lookup tables can be generated in parallel in $O(n^{1/2}\log n)$ work and $O(\log n)$ depth. The results for the sub-ranges should be represented using $O(\log\log n)$ bits each, and groups of $O(\log n/\log\log n)$ entries can be packed into a word as a post-processing step in $O(n/\log n)$ work and $O(\log n)$ depth.

Select. For select queries, we use Clark's select structure [5]. We describe the case for querying the position of 1 bits, and the case for querying 0 bits is done in the same way. Clark's data structure stores the position of every $\log n\log\log n$'th 1 bit, which defines *ranges*. For a range of length r between the positions, if $r \geq \log^2 n(\log\log n)^2$, then the select query answers are directly stored.

Otherwise, the position of every $\log r \log\log n$'th 1 bit is stored, which defines *sub-ranges*. For a sub-range of length r', if $r' \geq \log r' \log r (\log\log n)^2$ then answers are stored directly. Queries that fall into all other sub-ranges are answered via a lookup table that stores all answers for bit-strings of length $r' = O((\log\log n)^4)$. The following discussion assumes for simplicity that the range and sub-range boundaries coincide with half-word boundaries. This assumption can be removed by adding additional inputs to the lookup tables that indicate the portion of the half-word to be considered without affecting the overall complexity bounds.

To construct the select structure, we compute the number of 1's in each of the $2n/\log n$ half-words using table lookup, and perform a prefix sum over the $2n/\log n$ results. We can now identify all of the half-words that contain the position of a $k \log n \log\log n$'th 1 bit, for any integer k. For the ranges of length greater than $\log^2 n(\log\log n)^2$, we scan through the half-words in the range and store the location of every 1 bit. The location of all 1 bits within a half-word can be determined in $O(x)$ work and $O(1)$ depth via table lookup where x is the number of 1's in the half-word. The locations within the half-word are then offset by the location of the half-word in the sequence, again taking $O(x)$ work and $O(1)$ depth. Scanning the half-words takes $O(n/\log n)$ work and $O(\log n)$ depth. There are at most $n/(\log n \log\log n)$ locations of 1 bits found this way, and we can store their locations in the appropriate range in $O(n/(\log n \log\log n))$ work and $O(1)$ depth after a prefix sum (on a total of $O(n/\log n)$ half-words) that determines appropriate offsets.

For ranges of length less than $\log^2 n(\log\log n)^2$, we perform prefix sums over the half-words (as before, the count in a half-word is found via table lookup) in the range to identify which half-words have boundaries for sub-ranges, which takes $O(n/\log n)$ work and $O(\log n)$ depth overall. For the half-words that have boundaries, we output all of the boundary offsets in packed representation by using a table lookup. The lookup table takes a half-word, an offset j, and a length r, and outputs the location (relative to the beginning of the range) of every $j + k \log r \log\log n$'th 1 bit for all k in a packed representation. It can be constructed by considering all possible half-words, and all possible values of j and r, in $o(n/\log n)$ work and $O(\log n)$ depth. There are at most $O(n/(\log r \log\log n))$ boundaries, and each takes $O(\log\log n)$ bits to store. We can output $O(\log n/\log\log n)$ boundaries in a word in constant work, and so outputting all of the boundaries takes $O(n/(\log r \log n)) = O(n/\log n)$ work and $O(\log n)$ depth.

If answers in the sub-range need to be stored directly (i.e., the sub-range length r' is at least $\log r' \log r (\log\log n)^2$), then we store the answers relative to the start of the sub-range using $O(\log\log\log n)$ bits each. Directly generating the answers (as done in [22]) and then packing them into words would require $O(n/(\log\log n)^2)$ work since there could be that many answers, and this is too much. Therefore, we will generate the positions of all 1 bits relative to the start of the range in each half-word by using table lookup, where the result is packed into groups of $O(\log n/\log\log\log n)$ relative positions (each requires only $O(\log\log\log n)$ bits to represent). The lookup table also takes as input the number of positions to offset each answer. The offsets can be computed via a prefix sums over the counts of 1 bits in the half-words. The number of positions of 1 bits output is at most $O(n/\log\log^2 n)$, and so the number of groups is at most $O(n \log\log\log n/(\log n(\log\log n)^2)) = O(n/\log n)$. The last group in each half-word might not be fully packed but this adds at most $2n/\log n$ groups to the total. The offsets for storing the groups for each half-word can be pre-computed via prefix sums. The lookup table takes at most $\log^2 n(\log\log n)^2$ possible offsets, and has $2^{\log n/2} \log n$ entries per offset, so can be constructed in $o(n/\log n)$ work and $O(\log n)$ depth. The overall work for this step is thus $O(n/\log n)$ and the depth is $O(\log n)$. Finally, for the sub-ranges of length $r' < \log r' \log r (\log\log n)^2 = O((\log\log n)^4)$, the queries are answered via a look-up table that can be computed in $2^{O(\log\log n)^4} \cdot O((\log\log n)^4) = o(n/\log n)$ work and $O(\log n)$ depth.

For the select queries to work properly, all of the words inside each range and sub-range except the last should be fully packed, but this can be fixed with a post-processing step that generates

an array of new words, and computes for each old word where it should copy its results in the new word using a prefix sum. In parallel, each new word is then constructed sequentially from the corresponding old words. There are a total of $O(n/\log n)$ words in total, so this takes $O(n/\log n)$ work and $O(\log n)$ depth.

We have the following theorem for constructing rank/select structures on binary sequences:

Theorem 5.1. *The rank and select structures for a binary sequence of length n packed into $n/\log n$ words can be constructed in $O(n/\log n)$ work and $O(\log n)$ depth.*

5.2 Generalized Rank and Select Structures

In this section, we show how to construct rank and select structures on sequences with larger alphabets $\sigma = o(\log^{1/3} n)$ (this solution can also be used for binary sequences although the solution described in Section 5.1 is simpler). For a sequence of length n, Shun [22] describes how to construct the structures for $O(n)$ work and $O(\log n)$ depth. We show that the construction can be done in $O(n \log \sigma/\log n)$ work and $O(\log n)$ depth so that the rank/select structures can be used in the multiary wavelet tree described in Section 4. While a work bound of $O(n/\sqrt{\log n})$ suffices for use in the multiary wavelet tree, our goal is to match the work of the sequential algorithms for constructing the generalized rank/select structures of [1]. We assume the input is packed into $n \log \sigma/\log n$ words.

Rank. For the rank structure, a query $\mathsf{rank}_{\leq c}(S, i)$ returns the number of times a symbol *less than or equal to* c appears from positions 0 to i, which differs from the binary case. Thus simply creating σ copies of the binary rank structure, one for each character, will not suffice. We will instead use the generalized rank structure described in [1].

The generalized rank structure of [1] stores the σ ranks, one per character, for every $\sigma \log^2 n$'th symbol in the sequence. These symbols define ranges in the sequence, and we will refer to them as *range symbols*. For each range, the σ ranks of every $\log n/(3 \log \sigma)$'th symbol relative to the beginning of the range are stored. These symbols define sub-ranges, which we refer to as *sub-range symbols*. Queries inside a sub-range are of length at most $\log n/(3 \log \sigma)$ and can be answered in $O(1)$ work via table lookup. The table has $O(\sigma^{\log n/(3 \log \sigma)} \log n/(3 \log \sigma))$ entries per symbol, which sums to $o(n \log \sigma/\log n)$ overall, and thus can be constructed in $o(n \log \sigma/\log n)$ work and $O(\log n)$ depth using similar ideas as before.

We first describe how to compute the ranks of all sub-range symbols relative to the beginning of its sub-range. The algorithm requires pre-computing a lookup table which takes as input the current sub-range symbol, the previous sub-range symbol, and the block of $\log n/(3 \log \sigma)$ symbols in between, and outputs the generalized ranks of the current sub-range symbol relative to the beginning of the sub-range in $O(1)$ work (the result requires $\sigma \log(\sigma \log^2 n)$ bits and thus can fit in a word). The table has size $o(n \log \sigma/\log n)$ and can be constructed within the desired complexity bounds to support $O(1)$ work queries. Now the ranks of the sub-range symbols are computed in parallel using a prefix sums where the combining operator $\oplus$ is defined by the lookup table (note that the operator is associative, as required by prefix sums). The prefix sums will operate on entries e, which represents the combined ranks of one or more symbols that we are trying to compute ranks for. Each entry e stores the generalized ranks for the first and last symbol that it represents, which we denote by e_F and e_L, respectively. Now when combining an entry e with an entry e' that represents symbols earlier (to the left) in the input, we use the lookup table to update e_L given e'_L and the $\log n/(3 \log \sigma)$ symbols in between e'_L and e_F. We then form a new "combined" entry e'' that stores e'_F and e_L as the generalized ranks for its first and last symbols. This takes $O(1)$ work. The prefix sums computation uses a reduction tree, where initially each leaf is an entry representing one of the sub-range symbols. It then combines the entries in a binary fashion down

to a single entry representing the whole input. Finally, it propagates the information computed during the combining phase down the reduction tree (see [23] for details of prefix sums). At the end, the leaves of the tree will store the correct generalized ranks of the sub-range symbols. Since there are $3n \log \sigma / \log n$ symbols that we compute ranks for, the prefix sums takes $O(n \log \sigma / \log n)$ work and $O(\log n)$ depth. The results can be packed tightly into words containing using similar ideas as before.

To compute the generalized ranks for the range symbols, we perform a prefix sums over the first sub-range symbol of all ranges. When combining two entries, we can simply scan through all σ characters (in parallel) and update their generalized ranks. Each combining operation takes $O(\sigma)$ work and $O(1)$ depth, and there are $O(n/(\sigma \log^2 n))$ entries, giving a total complexity of $O(n/\log^2 n)$ work and $O(\log n)$ depth. The overall complexity for constructing the rank structure is $O(n \log \sigma / \log n)$ work and $O(\log n)$ depth.

Select. For the select structure, we could simply create σ copies of the binary select structure in Section 5.1, one per character. However, the select structure that we use takes $O(n/\log\log n)$ bits of space, and so this will not be a succinct representation for large σ. We will therefore parallelize the construction of the generalized select structure described in [1]. It has been described how to do this in $O(n)$ work in [22], but to do this in $O(n \log \sigma / \log n)$ work to match the bounds in [1] requires additional care.

In this structure, we have a separate select structure for each character but the structure is not the same as in the binary case. For a character c, the structure stores the location of every $\sigma \log^2 n$'th occurrence of c, and these occurrences define ranges (call these occurrences *range symbols*). For each range, if the length is at least $\sigma^2 \log^4 n$ then we store the answers directly, and otherwise we store the locations for every $\sigma(\log\log n)^2$'th occurrence of c relative to the start of the range, which define sub-ranges (call these occurrences *sub-range symbols*). For a sub-range, if the length is at least $\sigma^3(\log\log n)^4$, the answers are stored directly, and otherwise a lookup table is used to answer any query in the sub-range in $O(1)$ work. The table is of size $2^{\sigma^3(\log\log n)^4}\sigma^3(\log\log n)^4 = o(n \log \sigma / \log n)$ since $\sigma^3 = o(\log n)$. Thus it can be constructed within the desired complexity bounds. As before, the following discussion assumes that the the range and sub-range boundaries coincide with word boundaries, and this assumption can be removed without affecting the overall complexity bounds.

We split the input sequence into chunks of $\log^2 n/(3 \log \sigma)$ symbols and compute the number of occurrences of each character inside a chunk. For each group of $\log n/(3 \log \sigma)$ symbols in a chunk, we can output the number of occurrences of each character using table lookup. The table takes as input σ counts and a chunk, and outputs the updated counts based on the number of occurrences of each character in the input chunk. The table contains $O(2^{\log n/3}(\log n/(3 \log \sigma))\sigma \log\log n)$ entries, and thus can be computed in $o(n \log \sigma / \log n)$ work and $O(\log n)$ depth. Each answer from the table lookup uses $O(\sigma \log\log n)$ bits, which can be packed into a word and output in $O(1)$ work. We parallelize across chunks, where each chunk is processed sequentially, taking $O(\log n)$ depth since there are $\log n$ groups per chunk. The total work performed across all chunks is $O(n \log \sigma / \log n)$.

Now we must find the range symbols. We perform a prefix sums over the answers computed above, where the associative combining operator is defined by a lookup table that takes the σ counts from two chunks and outputs the σ counts that correspond to the sum of the counts from the two input chunks. The counts here will be relative to the beginning of the string, and thus an output can take $O(\sigma \log n)$ bits to represent and $O(\sigma)$ work to output. There are $O(n \log \sigma / \log^2 n)$ chunks, and thus the prefix sums takes $O(n\sigma \log \sigma / \log^2 n) = O(n \log \sigma / \log n)$ work and $O(\log n)$ depth.

We now know the number of occurrences of each character in each chunk as well as from the beginning of the string up to that chunk. This allows us to identify which chunks the range symbols occur in, and we search in the associated half-words in the chunk for the location of the range

symbol. For each chunk, we scan over the half-words sequentially updating the number of times we have seen a symbol so far via table lookup. Whenever we find a half-word that contains a range symbol, we use table lookup find the position of the j'th occurrence of a symbol inside the half-word in $O(1)$ work for an appropriate value of j. Thus, processing each chunk takes $O(\log n)$ work and depth. The lookup tables can be constructed in $o(n \log \sigma / \log n)$ work and $O(\log n)$ depth. This process gives all of the range symbols for a single character. There are at most $n/(\sigma \log^2 n)$ chunks that need to be checked per character, each one taking $O(\log n)$ work. Summed across all characters, the work is $O(n/\log n)$ and the depth is $O(\log n)$ (we can do this process for all characters and all chunks in parallel).

With this information, we can compute the lengths of the ranges between range symbols. For the ranges that are at least $\sigma^2 \log^4 n$ long, we store all locations of the character c which takes $n/(\sigma^2 \log^4 n) \cdot \sigma \log^2 n = O(n/(\sigma \log^2 n))$ work. Summed over all characters gives $O(n/\log^2 n)$ work. Finding these characters requires scanning the relevant chunks, which takes $O(n \log \sigma / \log n)$ work and $O(\log n)$ depth (recall that each chunk is scanned sequentially). If we mark the relevant chunks for each character beforehand, one scan over all of the chunks suffices to obtain the information for all characters, so this takes $O(n \log \sigma / \log n)$ work and $O(\log n)$ overall. We can pre-allocate space for these long ranges using prefix sums in the same amount of work and depth.

For ranges of length less than $\sigma^2 \log^4 n$, we compute the sub-range symbols. This process is mostly similar to how the range symbols were computed but since there can be up to $n/(\sigma(\log \log n)^2)$ sub-range symbols per character, outputting their positions directly would take too much work. We modify the lookup table to take as input offsets $j_1, \ldots, j_\sigma$, a character set C, and output the offset positions of every $j_c + k\sigma(\log \log n)^2$'th occurrence of $c \in C$ in a half-word for all integers k (j_c is set to -1 for $c \notin C$). The output positions are tagged with the corresponding character and given in a packed representation ($O(\log n / \log \log n)$ positions per word). The total work for writing out the positions of sub-range symbols will then be $O(n/(\sigma(\log \log n)^2)) \cdot \sigma \cdot \log \log n / \log n = o(n \log \sigma / \log n)$. We first mark the half-words with the characters that need to be considered, and then we scan over the chunks for all characters simultaneously using the lookup table described above. The overall work is $O(n \log \sigma / \log n)$ and depth is $O(\log n)$.

To determine sub-ranges of length at least $\sigma^3(\log \log n)^4$, we first return all sub-range starting locations for all characters that satisfy this inside a quarter-word ($\log n/4$ bits) using a lookup table. The table will return the (packed) positions of the sub-ranges that satisfy this property. The number of entries in the table is $O(2^{\log n/4})$ and can be constructed within our complexity bounds. However, since sub-ranges can span multiple quarter-words, we then use a prefix sum across all quarter-words where the associative operator is a lookup table that combines two quarter-words by keeping the latest position of each character in the first quarter-word and earliest position of each character in the second. It also outputs any sub-range starting locations where the difference between the latest position in the first quarter-word and the earliest position in the second is at least $\sigma^3(\log \log n)^4$. This table has $O(2^{\log n/2})$ entries and can again be constructed within the desired bounds. Since the output of the table lookups require $O(\sigma \log \log n)$ bits to store, it requires $O(1)$ to write out. The prefix sums across all quarter-words takes $O(n/\log n)$ work and $O(\log n)$ depth.

Finally, storing all of the occurrences for these sub-ranges uses a lookup table that takes a half-word and a list of characters that we want to output, and outputs the position (relative to the beginning of the sub-range) of every occurrence of an input character (tagged with the character) in a packed representation. Constructing this table can be done within the desired bounds. Performing the lookups across all half-words of large sub-ranges can be done in $O(n/\log n)$ work and $O(\log n)$ depth. There are at most $n/(\sigma^3(\log \log n)^4)$ sub-ranges requiring answers to be stored directly, each containing $\sigma(\log \log n)^2$ positions that require $O(\log \log n)$ bits each, and so outputting all such positions takes $O(n/(\sigma^2 \log n \log \log n))$ work.

Overall, constructing the generalized select structure takes $O(n \log \sigma / \log n)$ work and $O(\log n)$ depth, giving the following theorem:

Theorem 5.2. *For a sequence of length n containing characters in $[0, \ldots, \sigma - 1]$ packed into $n \log \sigma / \log n$ words, where $\sigma = o(\log^{1/3} n)$, the corresponding generalized rank and select structures can be constructed in $O(n \log \sigma / \log n)$ work and $O(\log n)$ depth.*

6 Conclusion

We have described parallel algorithms for wavelet tree construction with improved work complexity. The ideas extend to constructing wavelet trees of arbitrary shape as well as wavelet matrices. We also showed that the rank and select structures stored on the nodes of the wavelet tree can be constructed in the same work and depth bounds. An open question is to obtain a parallel wavelet tree algorithm with $O(n \log \sigma / \sqrt{\log n})$ work and polylogarithmic depth for any value of σ. It would also be interesting to implement this algorithm and experimentally compare its performance with existing parallel wavelet tree algorithms [11, 22, 15].

References

[1] M. A. Babenko, P. Gawrychowski, T. Kociumaka, and T. A. Starikovskaya. Wavelet trees meet suffix trees. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 572–591, 2015.

[2] P. C. P. Bhatt, K. Diks, T. Hagerup, V. C. Prasad, T. Radzik, and S. Saxena. Improved deterministic parallel integer sorting. *Information and Computation*, 94(1):29 – 47, 1991.

[3] R. D. Blumofe and C. E. Leiserson. Scheduling multithreaded computations by work stealing. *J. ACM*, 46(5):720–748, Sept. 1999.

[4] R. P. Brent. The parallel evaluation of general arithmetic expressions. *J. ACM (JACM)*, 21(2):201–206, 1974.

[5] D. R. Clark. *Compact Pat Trees.* PhD thesis, 1996.

[6] F. Claude and G. Navarro. Practical rank/select queries over arbitrary sequences. In *String Processing and Information Retrieval (SPIRE)*, pages 176–187, 2008.

[7] F. Claude and G. Navarro. The wavelet matrix. In *String Processing and Information Retrieval (SPIRE)*, pages 167–179. 2012.

[8] J. A. Edwards and U. Vishkin. Parallel algorithms for Burrows-Wheeler compression and decompression. *Theor. Comput. Sci.*, 525:10–22, Mar. 2014.

[9] P. Ferragina, G. Manzini, V. Mäkinen, and G. Navarro. Compressed representations of sequences and full-text indexes. *ACM Trans. Algorithms*, 3(2), May 2007.

[10] L. Foschini, R. Grossi, A. Gupta, and J. S. Vitter. When indexing equals compression: Experiments with compressing suffix arrays and applications. *ACM Trans. Algorithms*, 2(4):611–639, Oct. 2006.

[11] J. Fuentes-Sepulveda, E. Elejalde, L. Ferres, and D. Seco. Efficient wavelet tree construction and querying for multicore architectures. In *Symposium on Experimental Algorithms (SEA)*, pages 150–161, 2014.

[12] R. Grossi, A. Gupta, and J. S. Vitter. High-order entropy-compressed text indexes. In *ACM Symposium on Discrete Algorithms (SODA)*, pages 841–850, 2003.

[13] G. J. Jacobson. *Succinct Static Data Structures*. PhD thesis, 1988.

[14] R. M. Karp and V. Ramachandran. Parallel algorithms for shared-memory machines. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity (A)*, pages 869–942. 1990.

[15] J. Labeit, J. Shun, and G. E. Blelloch. Parallel lightweight wavelet tree, suffix array and FM-index construction. In *IEEE Data Compression Conference (DCC)*, 2016.

[16] V. Makinen and G. Navarro. Rank and select revisited and extended. *Theor. Comput. Sci.*, 387(3):332–347, 2007.

[17] C. Makris. Wavelet trees: A survey. *Comput. Sci. Inf. Syst.*, 9(2):585–625, 2012.

[18] J. I. Munro, Y. Nekrich, and J. S. Vitter. Fast construction of wavelet trees. In *String Processing and Information Retrieval (SPIRE)*, 2014.

[19] G. Navarro. Wavelet trees for all. In *Combinatorial Pattern Matching (CPM)*, pages 2–26. 2012.

[20] S. Rajasekaran and J. H. Reif. Optimal and sublogarithmic time randomized parallel sorting algorithms. *SIAM J. Comput.*, 18(3):594–607, 1989.

[21] R. Raman. The power of collision: Randomized parallel algorithms for chaining and integer sorting. In *Foundations of Software Technology and Theoretical Computer Science*, pages 161–175, 1990.

[22] J. Shun. Parallel wavelet tree construction. In *IEEE Data Compression Conference (DCC)*, pages 63–72, 2015.

[23] U. Vishkin. Thinking in parallel: Some basic data-parallel algorithms and techniques, 2010.